

## AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

### Listing of Claims:

- 1           1.       (Currently amended) A processor-implemented method for  
2       enabling efficient communication between a first and a second node in a network  
3       of by-routing network traffic through fault zones in the network, the method  
4       comprising:  
5           identifying a path from ~~a~~ the first network node to ~~a~~ the second network  
6       node;  
7           identifying a set of fault zones through which the identified path passes;  
8           for each fault zone in the set of fault zones, assigning as a zone weight the  
9       number of paths from the first network node to the second network node that  
10       include said fault zone, wherein the zone weight is determined from the path  
11       configuration of the network;  
12           calculating a path weight for the identified path, wherein said path weight  
13       is equal to the sums of said zone weights for each fault zone included in the  
14       identified path; and  
15           selecting the identified path as the current path for network traffic from the  
16       first node to the second node,  
17       whereby efficient communication from the first node to the second node is  
18       enabled along the selected path.
- 1           2.       (Original) The method of claim 1, further comprising:  
2           identifying a new path from the first network node to the second network

3 node;  
4 assigning zone weights to each fault zone in the new path;  
5 calculating a new path weight for the new path; and  
6 if the new path weight is lower than said path weight for the identified  
7 path, selecting the new path as the current path for network traffic from the first  
8 node to the second node.

1 3. (Currently amended) The method of claim 1, wherein:  
2 the first network node is identified by a first identifier;  
3 the second network node is identified by multiple identifiers, including a  
4 second identifier;  
5 selecting the identified path as the current path for network traffic from the  
6 first node to the second node comprises selecting the identified path as the current  
7 path for network traffic from the first ~~identifier~~ node to the second node, wherein  
8 the second node is identified by the second identifier; and  
9 paths other than the identified path are selected as the current paths for  
10 network traffic from the first node identifier to the second node, wherein the  
11 second node is identified by multiple identifiers other than the second identifier.

1 4. (Currently amended) A computer readable medium storing  
2 instructions that, when executed by a computer, cause the computer to perform a  
3 method for enabling efficient communication between a first and a second node in  
4 a network of by routing network traffic through fault zones in the network, the  
5 method comprising:  
6 identifying a path from a the first network node to a the second network  
7 node;  
8 identifying a set of fault zones through which the identified path leads;  
9 for each fault zone in the set of fault zones, assigning as a zone weight the

10 number of paths from the first network node to the second network node that  
11 include said fault zone, wherein the zone weight is determined from the path  
12 configuration of the network;  
13 calculating a path weight for the identified path, wherein said path weight  
14 is equal to the sums of said zone weights for each fault zone included in the  
15 identified path; and  
16 selecting the identified path as the current path for network traffic from the  
17 first node to the second node,  
18 whereby efficient communication from the first node to the second node is  
19 enabled along the selected path.

1 5. (Currently amended) A processor-implemented method for  
2 enabling efficient communication between a first and a second node in a subnet of  
3 by determining routing between nodes in a the subnet, the method comprising:  
4 identifying multiple fault zones in the subnet, each fault zone comprising  
5 one or more components of the subnet;  
6 configuring a central subnet manager to manage routing between nodes in  
7 the subnet;  
8 identifying a set of paths from a first node having a first identifier to a  
9 second node having multiple identifiers, including a second identifier, wherein  
10 traffic is deliverable to the second node using any of the multiple identifiers;  
11 for each fault zone traversed by one or more of the paths, establishing a  
12 zone weight based on the number of paths from the first node to the second node  
13 that traverse said fault zone, wherein the zone weight is determined from the path  
14 configuration of the subnet;  
15 for each path in the set of paths, establishing a path weight from the sums  
16 of the zone weights for each fault zone traversed by said path; and  
17 for each of the multiple identifiers of the second node, selecting as the

18 current path from the first ~~identifier node~~ to the second node identified by each of  
19 ~~the said identifier~~ multiple identifiers, from said set of paths, the path having the  
20 best path weight,  
21 whereby efficient communication from the first node to the second node is  
22 enabled along the selected path.

1 6. (Currently amended) A computer readable medium storing  
2 instructions that, when executed by a computer, cause the computer to perform a  
3 method for enabling efficient communication between a first and a second node in  
4 a subnet ~~of by~~ determining routing between nodes in ~~a~~ the subnet, the method  
5 comprising:  
6 identifying multiple fault zones in the subnet, each fault zone comprising  
7 one or more components of the subnet;  
8 configuring a central subnet manager to manage routing between nodes in  
9 the subnet;  
10 identifying a set of paths from ~~a~~ the first node having a first identifier to ~~a~~  
11 the second node having multiple identifiers, including a second identifier, wherein  
12 traffic is deliverable to the second node using any of the multiple identifiers;  
13 for each fault zone traversed by one or more of the paths, establishing a  
14 zone weight based on the number of paths from the first node to the second node  
15 that traverse said fault zone, wherein the zone weight is determined from the path  
16 configuration of the subnet;  
17 for each path in the set of paths, establishing a path weight from the sums  
18 of the zone weights for each fault zone traversed by said path; and  
19 for each of the multiple identifiers of the second node, selecting as the  
20 current path from the first ~~identifier node~~ to the second node, wherein the second  
21 node is identified by each of the said identifier multiple identifiers, from said set  
22 of paths, the path having the best path weight,

23        whereby efficient communication from the first node to the second node is  
24        enabled along the selected path.

1            7-13    (Canceled).

1            14.    (Currently amended) A system for enabling efficient  
2        communication between nodes in a subnet by determining routing through a the  
3        subnet comprising multiple fault zones, comprising:  
4            a network node configured to interface a client computing device with the  
5        subnet, wherein each node is identifiable by one or more node identifiers; and  
6            a client computing device, comprising:  
7                a processor; and  
8                a subnet manager module configured to determine routing between  
9            a first node and a second node in the subnet, wherein the first node is  
10          addressable by a first identifier and the second node is addressable by  
11          multiple identifiers, including a second identifier;  
12          wherein said subnet manager determines routing between the first node  
13        and second node by:  
14                for each fault zone in the subnet traversed by a path from the first  
15                node to the second node, calculating a zone weight based on the number of  
16                paths from the first node to the second node that traverse said fault zone,  
17                wherein the zone weight is determined from the path configuration of the  
18                subnet;  
19                for each of the paths from the first node to the second node,  
20                calculating a path weight based on the sums of said zone weights for the  
21                fault zones traversed by said path; and  
22                selecting as the current path from the first node identifier to the second  
23        node identified by the second identifier, the path ~~from the first identifier to the~~

24 ~~second identifier having the best path weight,~~  
25 whereby efficient communication from the first node to the second node is  
26 enabled along the selected path.

1 15. (Original) The system of claim 14, wherein the client computing  
2 device further comprises:  
3 a memory configured to store path weights of current paths between  
4 multiple pairs of node identifiers.

1 16. (Original) The system of claim 14, wherein said memory is further  
2 configured to store, in association with each of the current paths, zone weights for  
3 fault zones traversed by the current path.

1 17. (Original) The system of claim 14, wherein said subnet manager is  
2 further configured to disseminate routing information to a plurality of nodes in the  
3 subnet, said routing information including said current path from the first  
4 identifier to the second identifier.

18-24 (Canceled).